arXiv:1412.7955v1 [cs.NE] 26 Dec 2014

Unsupervised Learning Through Prediction in a Model of Cortex

Christos H. Papadimitriou* Santosh S. Vempala†

December 30, 2014

Abstract

We propose a primitive called PJOIN, for "predictive join," which combines and extends the operations JOIN and LINK, which Valiant proposed as the basis of a computational theory of cortex. We show that PJOIN can be implemented in Valiant's model. We also show that, using PJOIN, certain reasonably complex learning and pattern matching tasks can be performed, in a way that involves phenomena which have been observed in cognition and the brain, namely memory-based prediction and downward traffic in the cortical hierarchy.

1 Introduction

Human infants can do some amazing things, and so can computers, but there seems to be almost no intersection or direct connection between these two spheres of accomplishment. In Computer Science we model computation through algorithms and running times, but such modeling quickly leads to intractability, even when applied to tasks that are very easy for humans. The algorithms we invent are clever, complex and sophisticated, and yet they work in fashions that seem completely incompatible with our understanding of the ways in which the brain must actually work — and this includes learning algorithms. Accelerating advances in neuroscience have expanded tremendously our understanding of the brain, its neurons and their synapses, mechanisms, and connections, and yet no overarching theory appears to be emerging of brain function and the genesis of the mind. As far as we know, and the spectacular successes of neural networks [5, 8] notwithstanding, no algorithm has been proposed which solves some nontrivial computational problem in a computational fashion and style that can be credibly claimed to reflect what is happening in the brain when the same problem is solved. This is the scientific context that is motivating the present work.

We believe that formal computational models and algorithms can play a valuable role in bridging these gaps. We have been inspired in this regard by the pioneering work of Les Valiant on *the neuroidal model*, a random directed graph whose nodes and edges are capable of some very local computation [11]. We have also been prompted and influenced by certain interesting facts and conjectures about the brain's operation, which have been emerging in recent years. First, even though the visual cortex, for example, processes sensory input in a hierarchical way, proceeding from areas lower in the hierarchy to higher areas (from retina to V1, from there to V2, to V4, to MT, etc.), there seem to be many neuronal connections, and presumably significant firing traffic, directed *downwards* in this hierarchy (see e.g., [7] and subsequent work). Second, there seems to be an emerging consensus among students of the brain that *prediction* may be an important mode of brain function [1, 4, 9]: For example, vision involves not just of passive hierarchical processing of

*UC Berkeley. christos@cs.berkeley.edu
†Georgia Tech. vempala@gatech.edu

visual input, but also active acquisition of new input via saccades (rapid forays to new parts of the visual field), whose purpose seems to be to verify predictions made by higher layers of the visual cortex; significantly, many eye muscles are controlled by low areas of the visual hierarchy. Third, the connections in the cortex seem to deviate from the standard G_{np} model of random graphs, in that reciprocal connections (closing cycles of length two) and transitive connections (connecting nodes between which there is a path of length two *ignoring directions*) are much more likely than G_{np} would predict [10]. We briefly discuss a new model of brain connectivity in Section 6, as well as ways in which such an assumption complements our argument.

Valiant models the cortex as a random directed graph of *neuroids* (abstract neuron-like elements) connected via directed edges called *synapses*. The algorithms running on this platform are *vicinal*, by which it is meant that they are local in a very strict sense: the only communication from other neuroids that a neuroid is allowed to receive is the sum of the action potentials of all currently firing neuroids with synapses to it. This is the only outside information a neuroid can use in deciding what to do next. Valiant goes on to posit that real-world objects or concepts can be represented by *sets of neuroids*. There is a tradition of using such sets of neurons, or equivalently sparse binary vectors, as the basic formal unit of brain function [6], but Valiant's is, in our reading, the most explicitly algorithmic such framework. If one assumes that the underlying directed graph is random in the sense of the G_{np} model of random graphs [2], and taking into account what we know about the number of neurons in the cortex, the number of synapses out of (or into) a neuron, and the number of firing neurons that can cause another to fire, Valiant argues carefully [11] that it is possible to choose some number r, perhaps between 10 and 100, such that sets of r neurons can usefully represent a concept. Such sets are called *items*, and they constitute the basic element of Valiant's theory: all r neuroids of an item firing (or perhaps some overwhelming majority thereof) is coterminal with the corresponding concept being "thought about." Valiant shows that such items can be combined through the basic operations of JOIN and LINK, creating compound new items from, respectively simply connecting, existing items. For example, if A and B are items (established sets of about r neuroids) then Valiant shows how to create, via a vicinal algorithm, a new item JOIN(A, B) which stands for the combination, or conjunction, of the two constituent items.

Our main contribution is the formulation and implementation of a new operation on items that we call *predictive join* or *PJOIN*. This operation PJOIN extends JOIN in that, if only one of the constituent elements of PJOIN(A, B) fires, say A fires, then the structure will *predict* B, and initiate a sequence of downstream events whose purpose is to check this prediction. We show (Theorem 1) that the PJOIN of two items — which may be, and typically will be, themselves PJOINs — can be created by a vicinal algorithm so that it functions according to certain precise specifications, capturing the discussion above.

We suspect that the power of cortical computation lies not so much in the variety and sophistication of the algorithms that are actually implemented in cortex — which might be very limited — but in the ways in which the cortex interacts with an appropriately interesting environment through sensors and actuators. In other words, cortical architecture may happen to be very well adapted to the kinds of environments mammals live in. Here we use our PJOIN operation — *and almost nothing else* — to design a vicinal algorithm for performing a basic task that would be required in any such interaction with an environment: unsupervised learning of a set of patterns. By *learning* we mean that, if a pattern is presented to a sensory device for long enough, then our vicinal algorithm will "memorize it," in that it will create a structure of items connected by PJOINs which represent this pattern. If many patterns are presented, one after the other and with repetitions, all these patterns will be memorized in a shared structure. If sometime in the future one of the already presented patterns is presented again, then with high probability the structure will

"recognize it," by which we mean that there will be one special item (set of neuroids) whose firing will signal the act of recognition of that particular item. Moreover, a pattern that has already been presented once needs to be presented for a considerably shorter time on subsequent presentations. While JOIN can be viewed as the *unsupervised memorization* of two items in the neuronal model, PJOIN allows unsupervised memorization of patterns of arbitrary length. As we note in Section 4, using JOIN directly for more than two items does not yield this desired consistency in behavior.

We simulated these algorithms and measured their performance on binary patterns, including the sizes of representations created and updates to them, the total traffic generated in each presentation, and the fraction of traffic that is downward. The results, for varying ranges of parameters, are presented in Section 5. They are consistent with the theoretical predictions. In particular, both memorization and recognition are accomplished in a rather small number of "steps" — firings of individual neuroids — typically in the low tens; this is an important desideratum, because we know that complex cognitive operations can be performed in quite a bit less than second, and thus in fewer than a hundred such steps.

Finally, in Section 6 we discuss several important issues surrounding our results, such as the nature of the rudimentary *algorithmic control* necessary for the running of these and other vicinal algorithms, a new enhanced model of random graphs that appears to agree better with recent insights about cortical connectivity, as well as important further steps for this research program.

2 Valiant's Computational Model of Cortex

The basic structure in Valiant's computational theory of cortex is an *item,* which represents in the brain a real-world object or concept. Each item is implemented as a set of neurons. Valiant considers two alternative representations, one in which these sets are disjoint and one in which they may overlap; here we shall work in the disjoint framework, but we see no reason why our work could not be extended to overlapping items. Once such a set of neurons has been defined, the simultaneous firing of these neurons (or perhaps of some "overwhelming majority" of them) represents a brain event involving this item.

Valiant defined two key operations involving items. The JOIN operation creates, given two items A and B, a new item $C = \mathrm{JOIN}(A, B)$. Once the new item has been created, every time both A and B fire, C will fire in the next instant. A second operation is LINK. Given two established items A and B, the execution of $\mathrm{LINK}(A, B)$ has the effect that, henceforth, whenever item A fires, B will fire next. Valiant showed that these two operations — both the creation of these capabilities, and their application — can be implemented by simple *vicinal algorithms.* These algorithms are implemented in a particular formal *neuroidal model* of computation, involving only local computation, and communication strictly through action potentials. It is argued in [12, 14] that vicinal algorithms realistically capture the capabilities of neurons. Furthermore, it is argued in [13, 15] that even cascading usage of these two operations (taking the JOIN of two JOINs, and then again, and so on) is up to a point a reasonably realistic hypothesis.

The number r of neurons representing an item depends on the numerical parameters of cortex, as well as its random nature. Let n be the total number of neurons, and let d be the typical *out-degree* of a neuron, that is, the number of other neurons with which the neuron synapses. In the human brain, n is believed to be roughly 10^{11} and d roughly $10^4 - 10^5$. Let k be a number such that, if k neighboring neurons of a particular neuron are firing, then these k action potentials are typically enough to cause the neuron to fire. Valiant argues convincingly in [13] that, given our current knowledge of the parameters n, d, and k for various organisms, and assuming that the connections between neurons are random in the sense of G_{np}, a range exists for the new parameter

r, such that r neurons can usefully represent an item (see also [3]). The preceding discussion of the parameters k and r is an argument about the numerical plausibility of the neuroidal model, and does not mean that all items are represented by the same exact number of neurons, or that synapses are required to have the same strength and all neurons the same number of incoming synapses at all times — as we shall see, a neuroid can even set its own threshold in this model, and a synapse its own strength. However, for such a particular vicinal algorithm to be plausible, the numerical range of these parameters should be reasonable; for example, the firing of a neuron caused by the firing of very few other neurons, or a single one (either because of a tiny threshold or a huge synaptic strength) would be unreasonable. In the algorithms of this paper we do allow the parameters r and k, and even the firing threshold, to vary by a small factor in certain exceptional cases. Another desideratum on vicinal algorithms is that, even though the model assumes discrete time steps of a global clock, one should not rely too much on this. In other words, such algorithms should not have very long chains of consecutive steps requiring strict synchronization between many neuroids. Only sequences of up to three or four synchronized steps are used here (in [11, 12, 14] only two synchronized steps were needed). Naturally, it is necessary to assume that all neuroids comprising an item can fire simultaneously, in response to other items firing.

Vicinal Algorithms

The substrate of vicinal algorithms is a directed graph whose vertices are *neuroids*, (idealized neurons) connected through directed edges called *synapses*. At every discrete time $t \geq 0$, a neuroid v_i is at a *state* (T_i^t, f_i^t, q_i^t) consisting of three components: a positive real number T_i^t called its *threshold*; a Boolean f_i^t denoting whether i is *firing* presently; and another memory q_i^t which can take finitely many values. In our algorithms, T will for the most part be kept constant. Similarly, every synapse e_{ji} from neuroid v_j to neuroid v_i has a state (w_{ji}^t, qq_{ji}^t) consisting of a real number w_{ji}^t called its *strength;* plus another finitary memory qq_{ji}^t. When the superscript t is understood from context, it will be omitted.

The operation of vicinal algorithms is strictly local: neuroids and synapses update their own state in discrete time steps. The only way that the operation of neuroid v_i can be influenced by other neuroids is through the quantity $W_i = \sum_{j:f_j=1} w_{ji}$, the sum total of incoming action potentials. At each time instant, the firing status of all neuroids is updated:

$$W_i^t \leftarrow \sum_{j:f_j^t=1} w_{ji}^t; \quad \text{if } W_i^t \geq T_i^t \text{ then } f_i^{t+1} \leftarrow 1,$$

and after that the remaining of the state of the neuroids, and the state of the synapses, are updated:

$$(T_i^{t+1}, q_i^{t+1}) \leftarrow \delta(T_i^t, q_i^t, f_i^{t+1}, W_i^t),$$

$$(w_{ji}^{t+1}, qq_{ji}^{t+1}) \leftarrow \lambda(q_i^t, qq_{ji}^t, f_j^t, f_i^{t+1}, W_i^t).$$

That is, at each step f_i is set first — this is tantamount to assuming that the action potentials of other neuroids are transmitted instantaneously — and then the remaining state components of the neuron's state, and also of its incoming synapses, are updated in an arbitrary way. Notice that the update functions δ and λ are *the same for all neuroids v_i and synapses e_{ji}.*

This is the repertoire of actions of vicinal algorithms. For the algorithms to happen, some rudimentary control is necessary, causing the synchronous firing of certain neuroids, for example in the initiation of the creation of JOIN(A, B) explained next. We discuss the plausible vicinal implementation of such control in Section 6.

In our descriptions of our vicinal algorithms, instead of presenting explicit functions δ and λ, we shall use steps manipulating the *total state* of each neuroid, by which we shall mean all local information, namely the memory and threshold of the neuroid, plus the memory and strengths of all of its incoming synapses.

As an example of a vicinal algorithm, we next explain how the JOIN and LINK operations are performed; our exposition follows the one in Valiant [15].

The JOIN and LINK Operations

Given two items A and B, each represented by r neuroids, the operation JOIN(A, B) can be implemented in two steps, as follows:

- The neuroids that will represent $C = \text{JOIN}(A, B)$ will be recruited from a population of neuroids such that there are expected (on the basis of the random graph properties of the neuroidal network) to exist at least k synapses going from each of A, B to at least r of these neuroids, for some desired parameters r, k. These neuroids are initialtly at a total state that we call CANDIDATE: not firing, all memories at some null state, the standard threshold T, and all strengths of incoming synapses equal to $\frac{T}{k}$.

- At the first step, all neuroids in A fire. If, at the end of this step, a candidate neuroid fires (and therefore it had at least k synapses coming from A), then its total state becomes one that we call POISED: all synapses that come from firing A neuroids have now strength $\frac{T^2}{2kW_i}$ (so that if all neuroids of A fire again, they will together achieve a sum of $\frac{T}{2}$; recall that W_i is the total strength of all synapses from A, and therefore $\frac{kW_i}{T}$ denotes the number of incoming synapses from A). All candidates that did not fire enter a DISMISSED total state with all incoming strengths zero: they will not be needed to form C.

- Then in the next step all neuroids of B fire. If a poised neuroid fires in response, then it enters an OPERATIONAL state, in which it is ready to partake in the JOIN(A, B) operation, should both A and B fire simultaneously in the future. These are the neurons which will henceforth comprise item $C = \text{JOIN}(A, B)$. All synapses from B have strength $\frac{T^2}{2kW_i}$ where W_i denotes the new sum of strengths, while the neuroids from A retain their strength (and thus, if now all neuroids in A and B fire now, all neuroids in C will fire). All poised neuroids that did not fire in this second step are now DISMISSED.

We can summarize the vicinal algorithm for the creation of JOIN(A, B) as follows:

Step 1 $A \Rightarrow$ fire; not A and not $B \Rightarrow$ CANDIDATE

Step 2 CANDIDATE and FIRED $\Rightarrow$ POISED; $B \Rightarrow$ fire

After Step 2 POISED and FIRED $\Rightarrow$ OPERATIONAL; POISED and not FIRED $\Rightarrow$ DISMISSED

The operation LINK(A, B) which, once executed, causes all neuroids of B to fire every time A fires, is implemented in a similar manner, by employing a large population of intermediate *relay* neuroids [15] that are shared by all linking pairs.

Step 1 $A \Rightarrow$ fire; $B \Rightarrow$ PREPARED;

Step 2 Some relay neurons fire as a result of A's firing at Step 1; the neurons of B will likely all fire as a result of this.

After Step 2 PREPARED and FIRED $\Rightarrow$ L-OPERATIONAL

Here by L-OPERATIONAL it is meant that all incoming synapses from relay neuroids which have fired have memory state L-OPERATIONAL and strength equal to T/k, where T is the threshold and k is a small enough integer so that we expect there to be k relay nodes which both fire upon A firing *and* synapse to all neurons of B.

Refraction

Neurons in the brain are known to typically enter, after firing, a *refractory period* during which new firing is inhibited. In our model, the firing of a neuroid can be followed by a refractory period of R steps. This can be achieved through R memory states, say by increasing the neuroid's threshold T by a factor of 2^{R-1}, and subsequently halving it at every step. We shall routinely assume that, unless otherwise stated, the firing of an item will be followed by one refractory step.

3 Predictive JOIN

We propose now an enhanced variation of JOIN, which we call *predictive JOIN*, or PJOIN(A, B). As before, every time A and B fire simultaneously, C will also fire. In addition, if only A fires, then C will enter a state in which it will "predict B," in which C will be ready to fire if just B fires. Similarly, a firing of B alone results in the prediction of A. The effect is that the firing of A and B now need not be simultaneous for C to fire.

But there is more. Having predicted B, C will initiate a process whereby its prediction is checked: Informally, C will "mobilize" a part of B. If B also happens to be the predictive JOIN of some items D and E, say (as it will be in the intended use of PJOIN explained in the next section), then parts of those items will be mobilized as well, and so on. A whole set of items reachable from B in this way (by what we call *downward firings)* will thus be searched to test the prediction. As we shall see in the next section, this variant of JOIN can be used to perform some elementary pattern learning tasks.

Implementation of PJOIN

The vicinal algorithm for creating PJOIN(A, B) is the following:

Create PJOIN(A, B)

1. First, create $C =$ JOIN(A, B), by executing the vicinal algorithm of the previous section.
2. Each neuroid in C participates in the next steps with probability half. Valiant [12] allows vicinal algorithms to execute randomized steps, but here something mush simpler is required, as the random bit needed by a neuroid could be a part of that neuroid, perhaps depending on some component of its identity, such as its layer in cortex. We shall think of the chosen half (in expectation) of the neuroids of item C as comprising a new item which we call C_P, for "predictive C." That is, the neurons of item C_P are a subset of those of item C. (Strictly speaking this is departure from the disjoint model of implementing items, but it is a very inessential one.) The neuroids of C that are not in C_P now enter a total state called OPERATIONAL, ready to receive firings from A and/or B as any neuroid of JOIN(A, B) would.

3. Suppose now that A and B are predictive joins as well, and therefore they have parts A_P and B_P (this is the context in which PJOIN will be used in the next section). In the next two steps we perform LINK(C_P, A_P) and LINK(C_P, B_P), linking the predictive part of C to the predictive parts of its constituent items; there is no reason why the two LINK operations cannot be carried out in parallel. In the discussion section we explain that this step can be simplified considerably, omitting the LINK operations, by taking into account certain known facts about the brain's connectivity, namely the *reciprocity* of synapses. After these LINK creations, the working synapses from the relay neuroids to the neuroids of A_P and B_P will be in a memory state PARENT identifying them as coming "from above." Such synapses will have a strength larger than the minimum required one (say, double), to help the neurons identify firings from a parent.

4. After this is done, C_P enters the total state P-OPERATIONAL, where all strengths of synapses from A and B to C_P (but *not* to the rest of C) are doubled. The effect is that the neuroids of C_P will fire if *either* A fires, *or* B fires, *or both*. This concludes the creation of PJOIN(A, B). Notice that it takes four steps (we discuss in the last section one possible way to simplify it to three steps). After the creation of the predictive join is complete, indeed, if only one of A or B fires, then C_P will fire, initiating a breadth-first search for the missing item.

To summarize, we show below the vicinal algorithm, which we call **Create** PJOIN(A, B):

Steps 1 and 2 $C =$ JOIN(A, B);

After Step 2 C_P consists of about half of the neurons in C. C and not $C_P \Rightarrow$ OPERATIONAL

Steps 3 and 4 LINK(C_P, A_P); LINK(C_P, B_P)

After Step 4 $A_P, B_P \Rightarrow$ L-OPERATIONAL synapses in PARENT memory state;
$C_P \Rightarrow$ P-OPERATIONAL (synapses from A, B double their strength).

The *operation* of PJOIN(A, B) is a more elaborate version of that of JOIN:

1. If both inputs A and B fire simultaneously, then PJOIN(A, B) will operate as an ordinary JOIN.

2. As soon as one of A and B fires — suppose that A fires — then C_P fires (because of the doubled synaptic strengths), an event that will cause B_P to fire downwards in the next step. Notice that A_P will not fire as a result of C_P firing, because it is in a refractory state. Also, after the firing of A, the neuroids in C double the strength of the synapses coming from B (that is, from neuroids that do not come from a parent and are refractory), and all of C enters a total state called PREDICTING B. Thus, if henceforth the predicted item B fires, all of C will fire.

3. C_P's firing will cause B_P to fire downwards in the next step; A_P does not fire as a result of C_P firing, because it is in a refractory state. After this, C_P enters a PASSIVE total state, in which it will ignore firings of the E_P part of any item $E =$ PJOIN(C, D) further above. Again, this is accomplished by setting the strength of these synapses (which are the ones in the PARENT memory state) to zero.

4. If henceforth the predicted item B fires, then C fires and all of A, B, C go back to the OPERATIONAL (P-OPERATIONAL for C_P) total state.

5. Finally, if one of C's parents (by which we mean any item of the form PJOIN(C, D)) fires "downwards," then the neurons in C_P will fire (we think of them as firing "downwards"), thus propagating the search for the predicted items.

The state diagram of the desired operation of PJOIN(A, B) is shown in Fig. 1.

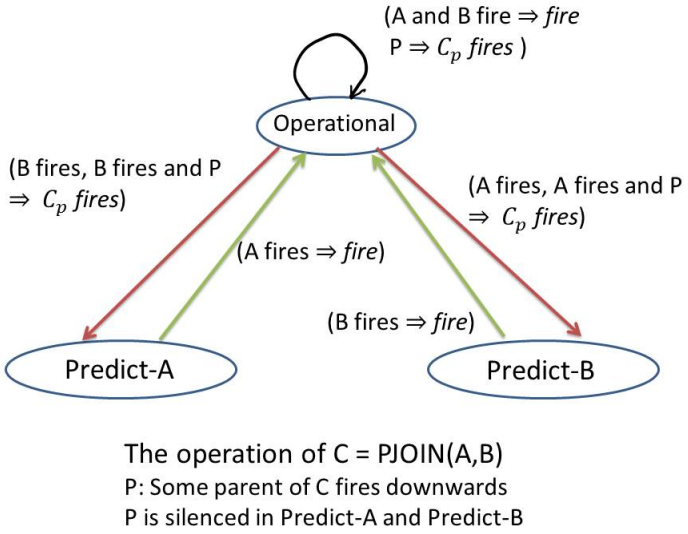
(A and B fire ⇒ fire
P ⇒ C_p fires)
Operational
(B fires, B fires and P
⇒ C_p fires)
(A fires, A fires and P
⇒ C_p fires)
(A fires ⇒ fire)
(B fires ⇒ fire)
Predict-A
Predict-B
The operation of C = PJOIN(A,B)
P: Some parent of C fires downwards
P is silenced in Predict-A and Predict-B

Figure 1: The *operational* state transition diagram for an item created by PJOIN

Theorem 1 *There is a vicinal algorithm that creates in four steps an item* PJOIN(A, B) *operating as specified by the diagram in Figure 1.*

Proof. We must argue two things: First, that the algorithm **Create** PJOIN(A, B) above is vicinal, that is, it can be implemented within the model described in Section 2 through appropriate state transformation functions δ and λ. And second, that the resulting item responds to firings of other items as specified by the diagram.

For the creation part, and since we assume that JOIN and LINK are implemented correctly, all we need to show is that the conditions that must fold after the second and after the fourth step of the algorithm **Create** PJOIN(A, B) can be implemented vicinally. The former is the identification of a random subset C_P of C, which is straightforward as discussed. The latter consists of (a) the PARENT memory state of the synapses from C_P to A_P and B_P, and the doubling of the strengths of these same synapses. Both (a) and (b) can be accomplished by the function λ of the neuroids right after the creation of LINK$(C_P.A_P)$ and LINK$(C_P.B_P)$.

To argue that the resulting PJOIN(A, B) operates as specified in the diagram, it sufficed to check each transition It is easy to see that the transition graph is complete). For example, for the transitions out of the OPERATIONAL state: The "A and B fire" self-loop follows from the fact that, in this case, PJOIN operates exactly as JOIN. The "P" self-loop is also immediate, since, for each parent D of the item, the synapses coming from D_P will cause the item to fire. And so on for the other transitions. □

Remark: Control and Synchrony. Vicinal algorithms should run with minimum control and synchrony. By *control* we mean firings initiated by outside events, other than the firing of neuroids

explicitly mentioned in the algorithm. The only control needed in our vicinal algorithms so far is the initiation of the firings of A and B, in consecutive steps, during the creation of JOIN(A, B). In Section 6 we shall discuss briefly how all control necessary for vicinal algorithms can actually be implemented vicinally.

Similarly, expecting billions of neurons to operate in precise lockstep is not realistic, and so vicinal algorithms should not require much synchrony. The only synchronization needed in our algorithms, besides the synchronous firing of all neurons of an item in response to the neurons of another item firing (a necessary assumption for Valiant's model, and whose plausibility we discuss in Section 6), is between the firing of two items in the two steps of the creation of JOIN(A, B) and the two steps of LINK(A, B) (and we shall soon argue that the latter one may not be necessary).

4 Pattern Memorization and Recognition

Cortical activity acquires meaning only when it interacts with the world via appropriate sensors and actuators. In this section we show that the primitive PJOIN enables a neurally plausible algorithm for memorizing and recognizing patterns.

Sensing the environment. By *pattern* we mean a binary vector with n coordinates, that is, a member of $\{0, 1\}^n$ for some parameter n. Visual patterns are only one of the many possibilities. In particular, we shall assume that there are n special items $S_1, \ldots, S_n$ called *sensors*. Each S_i has two special memory states, ZERO and ONE, so that the memory states of all n sensors "spell" the pattern.

Associated with each sensor S_i there are two other items which we denote by $\mathbf{0}_i$ and $\mathbf{1}_i$, such that:

- The states of the n sensors will remain unchanged during a period that we call *pattern presentation*. During such presentation, each sensor fires at most one — either spontaneously at a random time, or when predicted.
- The only activity of the neuroids of sensor S_i is the aforementioned firing, as well as the occasional change in memory state when a different pattern is presented.

The memory states of the n sensors typically remain unchanged from one time step to the next, displaying a pattern $x \in \{0, 1\}^n$. A maximum time period during which these memory states are unchanged is called a *presentation* of the corresponding pattern. We assume that, during such a presentation, the n sensor items fire spontaneously at random, each at most once. After the presentation of one pattern concludes, the presentation of a different pattern begins, and so on.

Memorization and recognition. We say that pattern $x \in \{0, 1\}^n$ has been *memorized* if, during its first presentation, a hierarchical structure of items is eventually formed with one top item $I(x)$ (that is, an item that does not participate in further PJOINs) which we think of as "representing" the presented pattern, in that, if x and y are different patterns, $I(x) \neq I(y)$. Notice that x's first presentation may follow the presentation of many other patterns $y, z, w \ldots$, and so its memorization structure will be "built" on top of existing memorization structures of the other patterns. Now if, in any subsequent presentation of the same pattern x, this special item $I(x)$ fires and is the highest-level item that fires, we say that the pattern has been *recognized*.

Learning with PJOIN. Before describing the algorithm for learning, we make a few remarks on sensor items.

- If sensor S_i fires in state ONE, item $\mathbf{1}_i$ fires next; similarly, if it fires in state ZERO,, $\mathbf{0}_i$ fires. This can be achieved, for example, by LINKing S_i to both $\mathbf{0}_i$ and $\mathbf{1}_i$, and appropriately manipulating the strengths of the corresponding synapses according to the memory state (ZERO or ONE) of S_i.

- The items $\mathbf{0}_i$ and $\mathbf{1}_i$ participate in PJOINs with other items, and thus they form the"basis" of the PJOIN structure which will be eventually formed in response to the current pattern presentation. By "basis" we mean that these are the only items in the structure that are not themselves PJOINs of other items (see Figure 2).

- There is one difference between these basis items and the other PJOIN items in the structure: As we have seen, if $C = \text{PJOIN}(A, B)$ then its predictive part, C_P, synapses to A_P and B_P. In contrast, if $C = \mathbf{1}_i$ or $\mathbf{0}_i$ then C_P synapses to the corresponding S_i item. Thus, every time $\mathbf{1}_i$ or $\mathbf{0}_i$ "fires downwards" S_i fires next, completing the chain of predictions. Importantly, it does so without refraction (because it may have to fire in the next step to verify the prediction).

- If and when a sensor item is predicted by a parent PJOIN, its sampling weight is increased, multiplied by a factor $M > 1$.

Memorization and recognition of a presented pattern is achieved by the algorithm shown in Figure 2. We say that an item is PJOIN-*eligible* if at least D time steps have elapsed after it has fired, and none of its parents has fired for the past D steps, where $D > 0$ is a parameter. Define the *level* of an item to be zero if out is an item of the form $\mathbf{0}_i$ or $\mathbf{1}_i$; and otherwise, if it is $\text{PJOIN}(A, B)$ to be one plus (recursively) the largest of the levels of A and B. In our algorithm, we allow the delay D associated with an item to depend on the level of the item. The selection of pairs to be PJOINed is as follows: PJOIN-eligible items are sampled, each with probability q. Note that, even though the algorithm of Figure 2 is depicted as running synchronously, the only synchrony that is necessary for its correctness is between the neuroids of the same item, and between pairs of items being joined. Furthermore, the only control (firings of neuroids initiated by events outside the algorithm, that is, by events other than firings of other neuroids), is the spontaneous firing of sensor items and the selection of PJOIN-eligible pairs of items to be PJOINed, and the initiation of the creation of the PJOIN. We shall discuss in the next section how such rudimentary control can be accomplished neuroidally.

We now turn to establishing the correctness of the algorithm. To state our main result of this section, we first review the pertinent parameters. p is the sampling probability of the sensors, and q the probability of sampling a PJOIN-eligible item for PJOINing. D is the delay after firing for an item to become PJOIN-eligible; we assume that $D = 2\ell + 2$. Patterns $x_1, \ldots, x_m \in \{0,1\}^n$ are presented with repetitions in some arbitrary order, and each presentation lasts at least T steps, where $T \geq 4(\log n)/p$.

Theorem 2 *Suppose that patterns $x_1, \ldots, x_m \in \{0,1\}^n$ are presented with repetitions in some arbitrary order, with each presentation lasting at least $T \geq 4\log n + (2\log n)/p$ steps; during each presentation, sensors are sampled with probability p, while eligible pairs are sampled with probability q; and the eligibility delay D satisfies $D \geq 2\ell + 2$ for items at level ℓ. Then, with high probability,*

- *At the first presentation of each pattern x_i, one particular highest-level item $I(x_i)$ is created, and during all subsequent presentations of x_i this item fires.*

LEARN(pattern x):

Initialization: set the memory state of item S_i to ZERO if $x_i = 0$ and to ONE if $x_i = 1$.

Repeat the following:

1. Sample sensor items which have not yet fired to fire next (with probability p).
2. Sample PJOIN-eligible pairs of items (A, B) to initiate the creation of $C = \text{PJOIN}(A, B)$ (with probability q).
3. Carry out one step of all neuroids.

until no neuroid fires or changes state.

Figure 2: Pattern memorization and recognition: sensor items and PJOIN-eligible pairs are sampled

- *Items $I(x_1), \ldots, I(x_n)$ are all distinct.*
- *Furthermore, after each pattern is presented $O(\log n + \log m)$ times, the $I(x_i)$ items do not participate in the creation of PJOINs.*

Remark 1. This consistent behavior, of essentially creating one PJOIN tree for each pattern (with trees having distinct roots but sharing subtrees), is not achievable using only JOIN. If mere JOINs were used, then on each presentation of a pattern of length $n > 2$, a new JOIN tree would be created, with no consolidation of prior knowledge. Another benefit of PJOIN is that, with it, the needed length of subsequent presentations of an item is decreased, since predictions are likely to accelerate considerably the rate with which sensors fire.

Remark 2. It is possible that, upon a presentation of an item which has already been memorized, new items may be created as PJOINs of old items. These new items should not necessarily be considered as undesirable "bugs;" they may be seen as alternative ways of parsing — of understanding — an existing pattern.

Turning to the proof of the theorem, we start with a bound on the maximum level of PJOIN items for a single pattern.

Lemma 1 *Under the conditions of the theorem, with high probability, the maximum level of an item is $O(\log n)$.*

Proof. Suppose n sensory items are PJOIN-eligible, so that $n - 1$ PJOINs happen. Order the PJOINs $1, 2, \ldots, n-1$. Fix a sensory item s and let X_j be the 0-1 random variable that indicates whether s participates in the j'th PJOIN. Then the height of s at the end of the process is $X = \sum_{j=1}^{n-1} X_j$. From this we have,

$$\mathsf{E}(X) \leq \sum_{j=1}^{n-1} \frac{1}{n-j} = H_{n-1} < 1 + \ln(n-1).$$

The claim follows from the Chernoff bound for independent Bernoulli random variables (with possibly different expectations):

$$\Pr(X - \mathsf{E}(X) > t\mathsf{E}(X)) \leq \left(\frac{e^t}{(1+t)^{(1+t)}} \right)^{\mathsf{E}(X)} \leq (e+1)^{-\mathsf{E}(X)}$$

where we used $t = e$. □

Proof of Theorem 2. First we note that any item at any level fires if and only if all its descendants fire in response to an input pattern. This is clear since an item fires if and only if both its children fire or it is an input level item activated by an input symbol. In the former case, this firing criterion recursively applied, leads to the set of all descendants of the original item, i.e., an item fires iff all its input level descendants fire. Thus, if the descendants of an item C are exactly the set of sensory items for a particular pattern x then C fires if and only if x is presented.

Next we claim that with high probability, when a new pattern x is encountered, the algorithm creates a new highest level cell to represent this pattern, i.e., the first presentation of the pattern x creates an item $I(x)$ whose descendants include all n sensory items that fire for x. After all possible processing of items according to the operational PJOIN state transition rules (see Fig. 1), if there is currently no item whose descendants match the pattern x, there must be a set of items S that are PJOIN-eligible. New items are then created by sampling a pair of items from S performing a PJOIN and adding the resulting item to S. This process finishes with a single item after $|S| - 1$ such PJOIN's. The descendants of this item correspond to the pattern x.

Now consider the presentation of a pattern x for which we already have a top-level item $I(x)$. All items whose descendants match the pattern will eventually fire. Thus, the original highest-level item for x will eventually fire, as claimed. En route, it is possible that two items that have fired, but none of whose parent items have fired as yet, will try to form a new PJOIN. First we bound the total number of rounds. Independent of the PJOIN creation, in each round, each sensor fires with probability p. Thus, with high probability, all n sensor items fire within $(2 \ln n)/p$ rounds. The additional number of rounds after all sensors fire is at most the depth of the tree, i.e., at most $4 \ln n$ whp. So the total number of rounds, again with high probability, is at most $4 \ln n + (2 \ln n)/p$. Thus our bound on T is adequate.

Next, on a presentation of x, consider any PJOIN item A with a parent B that is a descendant of $I(x)$. After item A fires, it will become PJOIN-eligible after D steps. But after A fires, its parent B predicts its sibling. The sibling is at level at most $\ell(A)$, and therefore it will fire within at most $2\ell(A) + 2$ steps. Therefore such an item A will not be PJOIN-eligible after D steps since its parent would have fired.

Now consider the point where all sensor items have been sampled and existing PJOIN items that will fire, including $I(x)$ have fired. There might be items that have fired but whose parents did not fire, items that are valid for x, but were created during the presentation of a different pattern y. These items can now form new PJOINs. However, after seeing each of the m input patterns at least once, no new PJOIN items will be created at level 1. This is because each sensory item, for each pattern, has at least one PJOIN parent. After another round of seeing all m input patterns, no PJOIN items are created at level 2. If we bound the maximum level of items, then this will bound the maximum number of rounds of seeing all m patterns after which no new PJOIN items are created.

In fact, the maximum level of an item will be $O(\log n + \log m)$. To see this, we we can either use Lemma 2 or the following more direct argument (with a slightly weaker bound) to conclude that the maximum number of items at level ℓ is at most $mn/2^{\ell}$. Fix a pattern i and suppose the first pattern and the i'th pattern share an r fraction of the sensor items. Then at level 1, some fraction of these rn sensors on which they agree will form PJOINs among themselves. The sensors involved in such PJOINs will no longer form new PJOINs on presentation of patterns 1 or i, since they already have PJOIN parents that will fire. The remaining common sensors could form new PJOINs that are valid for both 1 and i in later presentations of 1 or i. But no matter how many such PJOINs (among the common items) are formed on each presentation, the total number that can be valid for both 1 and i at level 1 is at most $rn/2$, since once a sensor participates in one such PJOIN (where the other sibling is also from the common set of r.n sensors), it will no longer do so.

For all m patterns, the number of PJOINs at level 1 valid for pattern 1 and any other pattern is at most $mn/2$.

We continue this argument upward. Of the items at level 1 (at most $rn/2$) that are valid for both 1 and i, on a presentation of i, some pairs could form new PJOINs that are valid for both 1 and i, if they do not already have such PJOIN parents with both children valid for both patterns, but no matter what order this happens, there can be at most $rn/4$ such PJOINs at level 2. This continues, with the number of PJOIN items halving at each level. Therefore, after $O(\log n + \log m)$ rounds of seeing all input patterns (in arbitrary order, and any positive number of times in each round), no further PJOINs are created and the top-level items are stable for each pattern. □

The next lemma gives a sharp bound on the number of PJOIN items valid for two different patterns.

Lemma 2 *Let $r_i = 1 - \frac{1}{n}H(x_1, x_i)$ be the fraction of overlap between two patterns x_1 and x_i, where H is the Hamming distance. The expected number of PJOINs created at level ℓ that are valid (i.e., will fire) for patterns x_1 and x_i is $(n/2)(r_i^4 + r_i^2(1-r_i^2)(1 + \frac{1}{(1+r)^2}))$ for $\ell = 1$ and*

$$\frac{n}{2^\ell} \cdot \left(r_i^{2^{\ell+1}} + r_i^{2^\ell}(1 - r_i^{2^\ell}) \cdot \left(1 + \left(\frac{1 + \left(\frac{1+(\ldots)^2}{1+r_i^{2^{\ell-1}}} \right)^2}{1 + r_i^{2^\ell}} \right)^2 \right) \right) < \frac{n}{2^\ell}.$$

for higher ℓ. Moreover, whp, the number of such PJOINs will be within a factor of 4 of this bound.

Proof. The proof of the expectation is by induction on ℓ. Consider two patterns x_1 and x_2 and let r be the fraction of sensors where they agree. Then on the first presentation of x_1, the number of PJOIN items created at level ℓ is $n/2^\ell$, and the expected fraction of them that are also valid for x_2 will be r^{2^ℓ} (since all constituent sensor items must be drawn from the common r fraction). Now, on a presentation of x_2, there might be more PJOINs created that are valid for x_1. At level 1, such PJOINs will come from sensors that are valid for both, but do not yet participate in PJOINs that are valid for both. The expected fraction of such sensors is $r - r^2 = r(1-r)$. For each such sensor, to be part of a new PJOIN valid for x_1, the other sensor must also be picked from this set. The total fraction of sensors participating in the creation of PJOINS will be $1 - r^2$, since an r^2 fraction of sensors already have PJOINs that will fire. Thus, the expected number of new PJOINs created at level 1 will be

$$\frac{n}{2} \cdot r(1-r) \cdot \frac{r(1-r)}{1-r^2} = \frac{n}{2} \cdot r^2(1-r^2) \cdot \frac{1}{(1+r)^2}$$

and thus the total number in expectation at level 1 is as claimed.

To compute the expected number of PJOINs created at the next level, we note that the first presentation of x_1 creates $r^4 \cdot (n/4)$ PJOINs that are valid for both x_1 and x_2. When x_2 is presented, the PJOIN items at level one that become PJOIN-eligible are $r^2(1-r^2)(1 + (1/(r+1)^2) \cdot (n/2)$, with the first term coming from PJOINs that were created on the presentation of x_1 and the second from new PJOINs created during the first presentation of x_2. Thus the number of PJOINs created at level 2 in expectation will be $n/4$ times

$$\frac{\left(r^2(1-r^2)\left(1 + \frac{1}{(1+r)^2}\right)\right)^2}{1-r^4} = r^4(1-r^4)\left(\frac{1 + \frac{1}{(1+r)^2}}{1+r^2}\right)^2.$$

Thus, the total number of PJOIN-eligible items at level 2 in expectation is

$$\frac{n}{4} \cdot \left(r^4 - r^8 + r^4(1-r^4) \left(\frac{1 + \frac{1}{(1+r)^2}}{1+r^2} \right)^2 \right) = \frac{n}{4} \left(r^4(1-r^4) f(2) \right).$$

Continuing, the expected number of PJOIN eligible items at level 3 is $(n/8)$ times

$$r^8(1-r^8) + \frac{\left(r^4(1-r^4)f(2)\right)^2}{1-r^8} = r^8(1-r^8) \left(1 + \left(\frac{f(2)}{1+r^4} \right)^2 \right) = r^8(1-r^8) f(3).$$

At the ℓ'th level, the expected number of PJOIN-eligible items is thus $(n/2^\ell)$ times

$$r^{2^\ell}(1-r^{2^\ell}) f(\ell) \quad \text{where } f(\ell) = 1 + \left(\frac{f(\ell-1)}{1+r^{2^{\ell-1}}} \right)^2.$$

Adding the number of PJOINs at level ℓ created by the first presentation of x_1 that are valid for both x_1 and x_2 and participated in PJOINs valid for both at the next level on the first presentation of x_1, which is $(n/2^\ell) r^{2^{\ell+1}}$, gives the claimed bound.

We proceed to show the inequality, namely that the above quantity, $r^{2^\ell}(1-r^{2^\ell}) f(\ell)$, is bounded by $1 - r^{2^{\ell+1}}$ for any $r \in [0,1]$, by induction on ℓ. In fact, our induction hypothesis is the following:

$$f(\ell) \leq 1 + \frac{1}{r^{2^\ell}}$$

It holds for $\ell = 2$. Assuming the induction hypothesis for $\ell - 1$,

$$\begin{aligned} f(\ell) &= 1 + \left(\frac{f(\ell-1)}{1+r^{2^{\ell-1}}} \right)^2 \\ &\leq 1 + \left(\frac{1 + \frac{1}{r^{2^{\ell-1}}}}{1+r^{2^{\ell-1}}} \right)^2 \\ &= 1 + \frac{1}{r^{2^\ell}}. \end{aligned}$$

Thus, $r^{2^\ell}(1-r^{2^\ell}) f(\ell) \leq (1-r^{2^\ell})(1+r^{2^\ell}) \leq 1 - r^{2^{\ell+1}}$.

The high probability bound follows from the independence of the PJOINs at each level. □

5 Experiments

We implemented PJOIN and tested it on binary input patterns. Besides checking for correctness, we measured the total number of items created (size of memory), as well as the total traffic and the number of downward predictions, for different values of pattern length n, number of patterns m, the number of inputs presentations T, the delay D, sensor item sampling probability p, and PJOIN-eligible item sampling probability q.

While we let each presentation run for as many steps as it took for every sensory item to be sampled, for 100-size patterns, no presentation lasted more than 50 steps and after a few presentations, this dropped to 15 or less steps. *Downward* traffic is indeed a substantial fraction of all traffic and increases with the number of distinct patterns presented. Figure 3(a) shows the the percent of downward traffic, as we increase the number of patterns m, for different values of pattern size

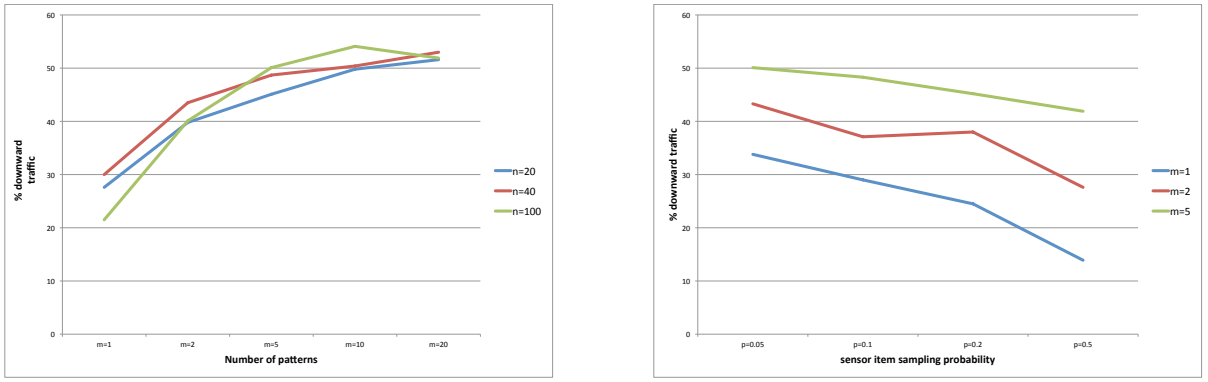

Figure 3: % of downward traffic with (a) more patterns (b) higher sensor sampling probability.

n, with $p = q = 0.1$ and $T = 100$ steps per presentation. Fig 3(b) is the downward traffic with higher sensor item sampling probability p, keeping $q = 0.1$ and $n = 40$. These results were obtained without any restriction on the level difference between two items being PJOINed.

Next, items created in early presentations are relatively stable, with few PJOINS created on later presentations, as predicted by Theorem 2. Figure 4 shows the average number of new PJOIN's created in each presentation, as the number of presentations increases, and as the delay for PJOIN eligibility (for items that have fired) is increased. For these experiments we kept n at 40 and $p = q = 0.1$.

Another experiment explores the sharing of common substructures by patterns. We generated the patterns by starting with one random pattern with $n = 100$, and obtain 9 more patterns by perturbing each coordinate of it randomly with probability p. In other words, the Hamming distance between consecutive patterns was p in expectation. We then used 10 presentations drawn randomly from this set of patterns. The perturbation probability p varied from 0.05 to 0.5 (the upper bound produces completely random patterns). The results in Figure 5 show that, as perhaps it could have been predicted, when the perturbation probability decreases (that is, when the patters are more similar), the total size of the structure created also decreases, as patterns share more substructures.

Finally, one may wonder whether unsupervised learning could be done more effectively by making prediction *random*. If a parent of an item $C = \text{PJOIN}(A, B)$ fires downwards, we currently require C_P to fire downwards to both A and B. One interesting variant would only fire downwards to *one* of A and B, chosen at random. This new version of PJOIN would generate considerably less downward traffic, and detect new patterns at a slower rate than the current one.

6 Some Remarks on Control

Vicinal algorithms are intended to be an austere, and therefore realistic, model of computation by the neurons in cortex. And yet certain aspects of the way we are using vicinal algorithms in this work require some further discussion and justification.

Perhaps the most striking aspect of the use of vicinal algorithms on items, as pioneered by Valiant, is the assumption that the r neuroids of an item can fire in synchrony, for example to initiate the creation of a JOIN or a PJOIN. We refer to this extraneous ability as "control." In

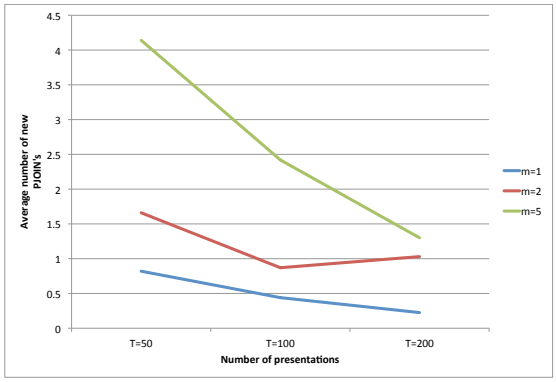

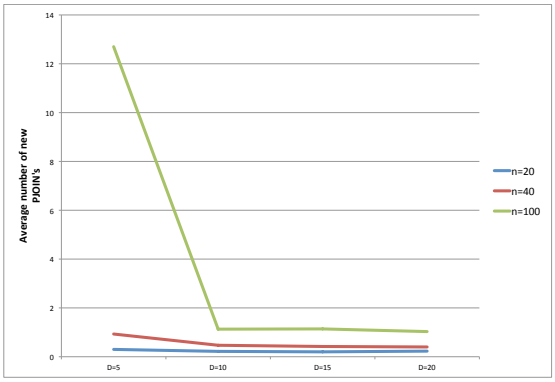

Figure 4: Average number of new PJOIN's per presentation, with (a) increasing number of presentations (b) higher delay for PJOIN eligibility.

the current section we briefly discuss how such control could be plausibly implemented within the vicinal model.

Reciprocity and a modified random graph model

Valiant's basic model assumes that neuroids are connected by a random directed graph of synapses in the G_{np} model in which all possible edges have the same probability p of being present. Measurements of the cortex, however, suggest a more involved model. In particular, Song et a.l [10] and subsequent work show that synaptic reciprocity (the chance that there is an edge from i to j *and from* j *to* i) is far greater than random. This is established for pairs chosen from quadruple whole-cell recordings, where the cells tend to be within close proximity of each other. The same is true of synaptic transitivity, i.e., the chance that two nodes connected by a path of length 2 are also connected directly.

A simple extension of the classical $G_{n,p}$ random graph model can account reasonably well for these departures from randomness. We propose the following: We start with n vertices. For every pair of vertices i, j within a distance threshold (in the case of the data from [10], this applies to every pair), with probability p only the $i \to j$ edge is in the graph; with probability p only $j \to i$; and with probability q *both* are in the graph; we assume that $2p + q < 1$. With the remaining $1 - (2p + q)$ probability, no arc exists between i and j.

This is already a useful model, incorporating the apparent reciprocity of cortical connections. A more elaborate *two-round model* also introduces transitivity. The first round is as above. In the second round, for each pair i, j with no arc between i and j, for every other proximal vertex k, if there is a path of length two between i and j from the first round, ignoring direction, then we add $i \to j$ with probability $r \cdot p/(2p + q)$, same for $j \to i$, and we add both arcs with probability $r \cdot q/(2p + q)$. This two-round model, with the specific values of $p = 0.05, q = 0.047$ and $r = 0.15$ gives a reasonable fit with the data in [10] (see Figure 6).

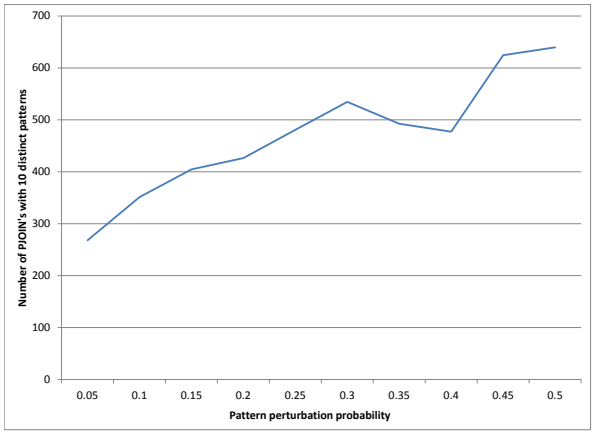

Figure 5: Number of PJOINs created, with decreasing overlap between patterns

An alternative implementation of PJOIN(A, B)

Once we assume that the neuroids have rich reciprocal connections between them, many possibilities arise. For a first example, the creation of PJOIN(A, B) can be simplified a little. After the neuroids of C have been chosen, C_P could consist of all neuroids in C which happen to have reciprocal connections back to A_P and B_P. Now the LINK step is unnecessary, and the creation of a PJOIN can be carried out in three steps, as opposed to four.

Using reciprocity for the control of items

Consider an item A. Is there a neuroid which is connected to each of the r neuroids through a path of reciprocal connections of the same small length, call it ℓ? For reasonable values of n (total number of available neuroids, say 10^9), r (number of neuroids in an item, say 50-100), and q (probability that two neuroids are connected in both directions, say 10^{-5}) it is not hard to see that such a neuroid is very likely to exist, with ℓ is bounded by

$$\frac{r-1}{r} \cdot \frac{\log n}{\log qn} \leq 3.$$

Furthermore, we can suppose that such a neuroid R_A ("the root of item A") is discovered by a vicinal process soon after the establishment of item A. One can then imagine that, once the neuroids of A fire, their root can sense this, and make them fire synchronously in a future step, to create a PJOIN(A, B), as required by our algorithm. The delay parameter D of our memorization algorithm may now be seen as the required number of steps for this synchronization to be set up.

Simultaneous PJOINs

Another important step of our memorization algorithm, whose implementation in cortex may seem problematic, is the selection of items to be PJOINed. How are they selected, and how are they matched in pairs? And, furthermore, won't the POISED neuroids (recall the creation of a JOIN described in Section 2) be "confused" by all other items firing at once? Valiant's algorithm for

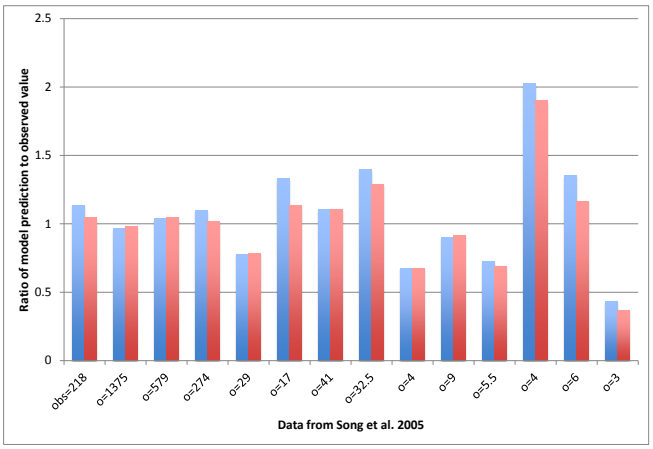

Figure 6: Comparison of the random model with reciprocity data. The first bar in each pair is with $p = q = 0.05$, the second with$p = 0.05, q = 0.047$.

creating JOINs assumes implicitly that nothing else is happening during these two steps; the simultaneous creation of several JOINs requires some thought and explanation.

Suppose that several PJOIN-eligible items fire, presumably in response to control signal from their respective roots, as hypothesized earlier. It is reasonable to assume that these firings are not simultaneous, but occur at times differing by a fraction of the duration of a step. One step after the first firing of an item A, and in response to that firing, several neuroids are POISED to become part of a PJOIN of A with another item. If we represent the firing of item A just by the symbol A, and the time of first availability of the POISED items for it as A', a sequence of events during our memorization algorithm can look like this:

$$\ldots A,\ B,\ C,\ A',\ D,\ B',\ C',\ E,\ F,\ D',\ G\ \ldots$$

where A' happens one step after A and similarly for B' after B, etc. Now this sequence of events will result in the following PJOINs being formed: PJOIN(A, D), PJOIN(B, E), PJOIN(C, E), PJOIN(D, G). Notice that E formed two PJOINs, with both B and C, and F formed none, and this is of course fine. This way, the spontaneous formation of PJOINs envisioned by our algorithm can happen without undesired "interference" between different PJOINing pairs, and with only minimal control.

We believe that this alternative way to select items to be PJOINed in our learning algorithm (by choosing *random stars* joining existing items, instead of random edges) enjoys the same favorable quantitative performance properties established by Theorem 2.

7 Conclusion and Further Work

We have introduced a new primitive, PJOIN, intended to capture the combining and predicting activity apparently taking place in cortex, and which can be implemented in Valiant's minimalistic model of vicinal algorithms. We showed that, by allowing items to spontaneously form PJOINs, starting from sensory inputs, complex patterns can be memorized and later recognized within

very reasonable time bounds. Much of the activity in this pattern recognition process consists of predicting unseen parts of the pattern, and is directed "downwards" in the hierarchy implied by PJOINs, in agreement with current understanding.

This work touches on the world's most challenging scientific problem, so there is certainly no dearth of momentous questions lying ahead. Here we only point to a few of them that we see as being close to the nature and spirit of this work.

- **Invariance.** We showed how PJOINs can be a part of the solution of a rather low-level problem in cognition, namely the memorization and recognition of patterns. Naturally, there are much more challenging cognitive tasks one could consider, and we discuss here a next step that is both natural and ambitious. One of the most impressive things our brain seems to accomplish is the creation of *invariants*, higher-level items capturing the fact that all the various views of the same object from different angles and in different contexts (but also, at an even higher level, all sounds associated with that object or person, or even all references to it in language, etc.), all refer to one and the same thing. It would be remarkable if there is simple and plausible machinery, analogous to PJOIN, that can accomplish this seemingly miraculous clustering operation. Incidentally, *language and grammar* seem to us to be important challenges that are related to the problem of invariance.

- **A PJOIN machine?** This work, besides introducing a useful cortical primitive, can be also seen as a particular stance on cortical computation: Can it be that the amazing accomplishments of the brain can be traced to some very simple and algorithmically unsophisticated primitives, which however take place at a huge scale?

- **Modeling the environment.** If indeed cortical computation relies not on sophisticated algorithms but on crude primitives like PJOIN, then the reasons for its success must be sought elsewhere, and most probably in two things: First, the *environment* within which the mammalian cortex has accomplished so much. And secondly, the rich *interfaces* with this environment, through sensors and actuators and lower-level neural circuits, resulting, among other things, in complex feedback loops, and eventually modifications of the environment. What seems to be needed here is a new theory of probabilistic models capable of characterizing the kinds of "organic" environments appearing around us, which can result — through interactions with simple machinery, and presumably through evolution — to an ascending spiral of complexity in cognitive behavior.

Acknowledgements. We are grateful to Tatiana Emmanouil and Sebastian Seung for useful references, and to Vitaly Feldman and Les Valiant for helpful feedback on an earlier version of this paper.

References

[1] Patricia S. Churchland, V. S. Ramachandran, and Terrence J. Sejnowski. A critique of pure vision. In Cristof Koch and Joel L. Davis, editors, *Large Scale Neuronal Theories of the Brain*. MIT Press, Cambridge, MA, 1994.

[2] Paul Erdos and Alfred Renyi. On the evolution of random graphs. *Publ. Math. Inst. Hungary. Acad. Sci.*, 5:17–61, 1960.

[3] Vitaly Feldman and Leslie G. Valiant. Experience-induced neural circuits that achieve high capacity. *Neural Computation*, 21(10):2715–2754, 2009.

[4] Jeff Hawkins and Sandra Blakeslee. *On Intelligence*. Times Books, 2004.

[5] G E Hinton and R R Salakhutdinov. Reducing the dimensionality of data with neural networks. *Science*, 313(5786):504–507, July 2006.

[6] P. Kanerva. *Sparse Distributed Memory*. A Bradford book. MIT Press, 1988.

[7] Victor A. F. Lamme and Pieter R. Roelfaema. *Trends in Neurosciences*, page 571.

[8] Yann LeCun, Lon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. In *Proceedings of the IEEE*, volume 86, pages 2278–2324, 1998.

[9] R.P.N. Rao and D.H. Ballard. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, 2:79–87, 1999.

[10] Song S, Sjöström PJ, Reigl M, Nelson S, and Chklovskii DB. Highly nonrandom features of synaptic connectivity in local cortical circuits. *PLoS Biology*, 3(3), 2005.

[11] Leslie G. Valiant. *Circuits of the mind*. Oxford University Press, 1995.

[12] Leslie G. Valiant. A neuroidal architecture for cognitive computation. *J. ACM*, 47(5):854–882, 2000.

[13] Leslie G. Valiant. Memorization and association on a realistic neural model. *Neural Computation*, 17(3):527–555, 2005.

[14] Leslie G. Valiant. A quantitative theory of neural computation. *Biological Cybernetics*, 95(3):205–211, 2006.

[15] Leslie G. Valiant. The hippocampus as a stable memory allocator for cortex. *Neural Computation*, 24(11):2873–2899, 2012.